

REMARKS

Claims 1, 11, and 21 have been amended.

Claims 1 – 30 are present in the subject application.

In the Office Action dated February 7, 2007, the Examiner has rejected claims 1 – 30 under 35 U.S.C. §103(a). Favorable reconsideration of the subject application is respectfully requested in view of the following remarks.

The Examiner has rejected claims 1 – 30 under 35 U.S.C. §103(a) as being unpatentable over U.S. Patent No. 4,956,774 (Shibamiya et al.) in view of U.S. Patent No. 6,006,220 (Haderle et al.). Briefly, the present invention is directed toward enhancing the access path selection process by storing the relevant frequency distribution and the host variable values used for selecting the access path together with the corresponding access path executable. This way, the host variables impact to the access path can be inexpensively assessed and the re-optimization triggered only if necessary.

In order to assist in an understanding of the invention, the Examiner's attention is respectfully directed to the example in the specification (e.g., Page 4, line 9 to Page 5, line 6). The example provides a table with three columns (C1, C2 and C3) and an index (C1, C2). A sample query is executed against this table including a host variable (hv) with successive values of 'B', 'A' and 'D' assigned thereto during successive executions of the sample query. The first execution of the sample query (hv = 'B') produces an access path and corresponding information is stored. At the next execution of the sample query (hv = 'A'), a check reveals that the value of 'A' for the host variable (hv) will not change the previous access path (e.g., since the cardinality of 'A' is close to the cardinality of the previous value 'B' in the example), therefore an access

path is not re-generated. However, at the subsequent execution of the sample query (hv = 'D'), the check indicates that the value of 'D' for the host variable (hv) will change the previous access path (e.g., since the cardinality of 'D' is not close to the cardinality of the previous value 'A' in the example). In this case, the access path is re-generated with value 'D' for the host variable.

The Examiner takes the position that the Shibamiya et al. patent discloses the claimed subject matter, except for generating the preferred access path for use with a subsequent execution of the same query based upon the information related to the second value received during the subsequent execution differing sufficiently from the stored information related to the first value received during the execution time. The Examiner further alleges that the Haderle et al. patent discloses this feature and that it would have been obvious to combine the Shibamiya et al. and Haderle et al. patents to attain the claimed invention.

Independent claims 1, 11 and 21 are considered to overcome the Shibamiya et al. and Haderle et al. patents for at least the claimed features of the query variable comparison to regenerate an access path. However, in order to expedite prosecution of the subject application, independent claims 1, 11 and 21 have been amended and recite the features of: generating a preferred access path for a query with at least one variable at execution time, wherein the preferred access path is generated based upon a first value for each variable in the query produced and received during the execution time; and regenerating the preferred access path for the query in response to the comparison indicating that the information related to the second value received during the subsequent execution differs sufficiently from the stored information

related to the first value received during the execution time to enable generation of an access path different than the preferred access path.

In other words, the independent claims recite that the first value for each query variable is produced and received at execution time which is utilized to generate a preferred access path during that execution time, and regenerating the access path for a subsequent execution of the same query with each query variable receiving a second value during the subsequent execution in response to sufficient differences between the first and second values for the query variables.

The Shibamiya et al. patent does not disclose, teach or suggest these features. Rather, the Shibamiya et al. patent discloses a method for more accurately estimating the time required to process a database query using a selected index. A preferred embodiment consists of two major phases. In the collection phase, frequency of occurrence statistics for the most frequent values are collected and saved for each of the database systems indexes in a system catalog. This collection is initially done after each table has been initially populated (loaded) with records and indexed. Since there are usually several possible alternative access paths for each query, the bind phase includes an optimizer to select the best access path for the query. In the second phase, the collected statistics saved in the system catalog are utilized by the optimizer during the bind process to estimate the time required to use each of the various access paths available to satisfy each query. The accurate estimates provided allow the optimizer to reliably select the best access path for each query (e.g., See Fig. 1; Column 4, lines 60 - 63; and Column 5, line 15 to Column 6, line 15).

Thus, the Shibamiya et al. patent discloses the optimizer being utilized to select an access path during the bind phase or process (e.g., prior to execution as shown in Fig. 1), as opposed to

Amendment

U.S. Patent Application Serial No. 10/688,951

generating a preferred access path at execution time as recited in the independent claims. Further, the Shibamiya et al. patent collects frequency of occurrence statistics to determine a time estimate for each access path associated with a query, and selects the access path based on the time estimates. There is no disclosure, teaching or suggestion of comparing information related to values of variables within different executions of the same query or, for that matter, regenerating the access path in response to sufficient differences between that information enabling generation of a different access path as recited in the independent claims.

The Haderle et al. patent does not compensate for the deficiencies of the Shibamiya et al. patent and similarly does not disclose, teach or suggest the above features. Rather, the Haderle et al. patent discloses the RDBMS selecting a default filter factor (construed by the Examiner as the claimed first value of the query variables) during the BIND process (e.g., or at compile time) to select an access path. The filter factor (also known as selectivity) is determined by the RDBMS and refers to a quantity of qualified rows of a table that will be satisfied for the query. Since the values of the query variables have unknown values at compile or BIND time, the default filter factor varies from the filter factor produced with actual values of the variables at execution time. Consequently, the access path selected by the optimizer utilizing the default filter factor selected by the RDBMS is not optimal (e.g., See Column 1, line 57 to Column 2, line 1; Column 3, lines 63 - 66; and Column 4, lines 8 - 10, 13 - 16, 21 - 24, 30 - 32 and 56 - 63). Accordingly, the selected access path is re-optimized at execution time when the query variables are known in order to provide a more accurate estimate of the filter factor for selecting the access path (e.g., See Column 4, line 56 to Column 5, line 5).

Amendment
U.S. Patent Application Serial No. 10/688,951

The Examiner takes the position that the Haderle et al. patent teaches two types of variables, namely default and actual. The Examiner further indicates that the default value (or default filter factor disclosed in the Haderle et al. patent) corresponds to the claimed first value and is received during execution time, while the actual value (or estimated filter factor disclosed in the Haderle et al. patent) corresponds to the claimed second value that differs sufficiently from the first value.

However, Applicant respectfully disagrees with the Examiner's interpretation of the Haderle et al. patent. Initially, the filter factor (also known as selectivity) of the Haderle et al. patent is determined by the RDBMS to select an access plan and refers to a quantity of qualified rows of a table that will be satisfied for the query. Thus, neither the default nor actual (estimated) filter factors are received by variables in a query as recited in the independent claims.

Further, the filter factor is determined by the RDBMS based on actual values of the query variables to select an access plan. Since the actual (estimated) filter factor cannot be determined at BIND or compile time due to the values for query variables being unknown, the RDBMS selects the default filter factor at BIND or compile time in order to initially select the access plan. Thus, the default filter factor is clearly selected at BIND or compile time, as opposed to the first value for the query variables being produced and received during execution of the query as recited in the independent claims.

Moreover, the Haderle et al. patent discloses that the filter factor is estimated at execution time (actual value) when values for the query variables are available to enable determination of the filter factor with enhanced accuracy and selection of a more optimal access path (e.g., See Column 4, line 56 to Column 5, line 5). Thus, the default filter factor is utilized at BIND or

Amendment

U.S. Patent Application Serial No. 10/688,951

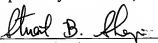
compile time of a query, while the estimated filter factor is determined at execution time of the query, as opposed to first and subsequent executions of the same query as recited in the independent claims. In addition, there is no disclosure, teaching or suggestion of comparing information related to values of variables within different executions of the same query or, for that matter, regenerating the access path in response to sufficient differences between that information enabling generation of a different access path as recited in the independent claims.

Since the Shibamiya et al. and Haderle et al. patents do not disclose, teach or suggest the features recited in independent claims 1, 11 and 21 as discussed above, these claims are considered to be in condition for allowance.

Claims 2 – 10, 12 – 20 and 22 - 30 depend, either directly or indirectly, from independent claims 1, 11 and 21, respectively, and therefore, include all the limitations of their parent claims. These dependent claims are considered to be in condition for allowance for substantially the same reasons discussed above in relation to their parent claims and for further limitations recited in the dependent claims.

The application, having been shown to overcome the issues raised in the Office Action, is considered to be in condition for allowance and a Notice of Allowance is earnestly solicited.

Respectfully submitted,


Stuart B. Shapiro
Registration No. 40,169

EDELL, SHAPIRO & FINNAN, LLC
1901 Research Blvd., Suite 400
Rockville, Maryland 20850-3164
(301) 424-3640
Delivered: 4/24/07